

WHAT IS CLAIMED IS:

1. In a system comprised of a plurality of objects where multiple copies of an object may exist, a method for maintaining consistent copies of the object, comprising the steps of:

5 applying a plurality of consistency policies in which application of at least one consistency policy results in different system performance than a second consistency policy; and

10 selecting a consistency policy from the plurality of consistency policies for an object, wherein the selection is made to improve system performance.

15 2. The method as recited in claim 1, wherein the at least one consistency policy includes an update-all consistency policy and the second consistency policy includes an update-holders consistency policy.

20 3. The method as recited in claim 1, wherein the at least one consistency policy includes a coordinate-all consistency policy and the second consistency policy includes a coordinate-holders consistency.

4. The method as recited in claim 1, further comprising including in the plurality of consistency policies strong and weak consistency policies.

5

5. The method as recited in claim 1, further comprising including in the plurality of consistency policies a strong consistency under at least one condition but a weak consistency policy if the at least one condition is not met.

10

6. The method as recited in claim 1, further comprising a step of managing the plurality of consistency policies using a consistency coordinator.

15

7. The method as recited in claim 1, wherein the step of selecting is performed by an application, which writes the object.

20

8. The method as recited in claim 1, wherein an object has a lifetime and the method further comprises a step of switching a consistency policy of the object during the object's lifetime.

9. The method as recited in claim 1, further comprising steps of:

measuring activity of a consistency coordinator, which manages the consistency policies in the system; and

5 maintaining connections with caches in the system in accordance with the activity of the consistency coordinator.

10. The method as recited in claim 9, further comprising communicating the activity of the consistency coordinators to the caches.

11. The method as recited in claim 10, wherein the step of communicating the activity comprises sending heartbeat messages to the caches.

15 12. The method as recited in claim 1, wherein the step of selecting includes choosing a consistency policy for at least one object, which maximizes system performance.

20 13. The method as recited in claim 12, wherein system performance is maximized by adjusting at least one of CPU overhead, communication latency and message overhead.

14. The method as recited in claim 1, wherein a consistency policy of at least one object is specified as a condition in terms of a temporal or semantic state of the object.

5

15. The method as recited in claim 1, wherein the consistency policy is selected from at least one of always strong consistency, conditional strong consistency, weak consistency with guarantees, and weak consistency.

10

16. The method as recited in claim 1, further comprising one of differentiating and prioritizing communication between a cache and a consistency coordinator by a cache device.

15

17. The method as recited in claim 16, further comprising maintaining at least two queues in the cache to hold messages communicated to the consistency coordinator.

20

18. The method as recited in claim 17, further comprising the step of prioritizing messages in one queue with a higher priority than messages in another queue.

19. The method as recited in claim 16, further comprising the step of maintaining a number of connections by a cache which is dynamically varied depending upon a load on the consistency coordinator.

5

20. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for a method for maintaining consistent copies of the object, the method steps comprising:

applying a plurality of consistency policies in which application of at least one consistency policy results in different system performance than a second consistency policy; and

15 selecting a consistency policy from the plurality of consistency policies for an object, wherein the selection is made to improve system performance.

21. In a system comprised of a plurality of objects in which multiple copies of an object may exist, a method for maintaining consistent copies of the object, comprising the steps of:

maintaining consistency using a plurality of consistency

policies in which at least one consistency policy achieves stronger consistency results than a second consistency policy; and

5 selectively choosing a consistency policy for at least one object, which balances between consistency level and performance.

22. The method as recited in claim 21, further comprising a step of adjusting a level of consistency for at 10 least one object in response to consistency overhead.

23. The method as recited in claim 21, wherein an object managed using one of expiration time, update all, update holders, and deferred invalidation consistency becomes 15 managed using strong consistency.

24. The method as recited in claim 21, wherein an object managed using strong consistency becomes managed using one of update all, update holders, and deferred invalidation 20 consistency.

25. The method as recited in claim 21, wherein the at least one consistency policy includes an update-all

consistency policy and the second consistency policy includes an update-holders consistency policy.

26. The method as recited in claim 21, wherein the at
5 least one consistency policy includes a coordinate-all
consistency policy and the second consistency policy includes
a coordinate-holders consistency policy.

27. The method as recited in claim 21, further
10 comprising including in the plurality of consistency
policies, strong and weak consistency policies.

28. The method as recited in claim 21, further
comprising including in the plurality of consistency
15 policies, a strong consistency under at least one condition
but a weak consistency policy if the at least one condition
is not met.

29. The method as recited in claim 21, further
20 comprising a step of managing the plurality of consistency
policies using a consistency coordinator.

30. The method as recited in claim 21, wherein the step of selectively choosing a consistency policy is performed by an application, which writes the object.

5

31. The method as recited in claim 21, wherein an object has a lifetime and the method further comprises the step of switching a consistency policy of the object during the object's lifetime.

10

32. The method as recited in claim 21, further comprising steps of:

measuring activity of a consistency coordinator, which manages the consistency policies in the system; and

15 maintaining connections with caches in the system in accordance with the activity of the consistency coordinator.

15

33. The method as recited in claim 32, further comprising communicating the activity of the consistency coordinators to the caches.

20

34. The method as recited in claim 32, wherein the step of communicating the activity includes sending heartbeat

messages to the caches.

35. The method as recited in claim 21, wherein the step of selectively choosing includes choosing a consistency policy for at least one object which maximizes system performance.

36. The method as recited in claim 35, wherein system performance is maximized by adjusting at least one of CPU overhead, communication latency and message overhead.

37. The method as recited in claim 21, wherein a consistency policy of at least one object is specified as a condition in terms of a temporal or semantic state of the object.

38. The method as recited in claim 21, wherein the consistency policy is selected from at least one of always strong consistency, conditional strong consistency, weak consistency with guarantees, and weak consistency.

39. The method as recited in claim 21, further comprising one of differentiating and prioritizing

communication between a cache and a consistency coordinator by a cache device.

40. The method as recited in claim 39, further
5 comprising maintaining at least two queues in the cache to hold messages communicated to the consistency coordinator.

41. The method as recited in claim 40, further comprising the step of prioritizing messages in one queue
10 with a higher priority than messages in another queue.

42. The method as recited in claim 39, further comprising the step of maintaining a number of connections by a cache which is dynamically varied depending upon a load
15 on the consistency coordinator.

43. A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for a method for
20 maintaining consistent copies of an object, the method steps comprising:

maintaining consistency using a plurality of consistency policies in which at least one consistency policy achieves

stronger consistency results than a second consistency policy; and

selectively choosing a consistency policy for at least one object, which balances between consistency level and 5 performance.

44. A system for maintaining consistent copies comprising:

a plurality of caches for storing objects wherein 10 multiple copies of an object may exist; each cache comprising at least two queues, which designate an update priority of the object included in that queue;

a plurality of consistency policies maintained 15 throughout the system such that at least one consistency policy results in different performance than a second consistency policy; and

a coordination coordinator having selective 20 communication with the caches, which manages requests for updates from the caches in accordance with the queue priority.

45. The system as recited in claim 44, wherein one

consistency policy includes update-all consistency and a second policy includes update-holders consistency.

46. The system as recited in claim 44, wherein one consistency policy includes coordinate-all consistency and a second policy includes coordinate-holders consistency.

47. The system as recited in claim 44, wherein the plurality of consistency policies includes strong and weak consistency policies.

48. The system as recited in claim 44, wherein the plurality of consistency policies includes a strong consistency for an object under at least one condition but a weak consistency policy for the object if the at least one condition is not met.

49. The system as recited in claim 44, further comprising an application, which writes the object, for selecting the consistency policy for an object.

50. The system as recited in claim 44, further comprising a number of connections between the consistency

coordinator and the caches wherein the number is adjusted in accordance with activity of the consistency coordinator.

5 51. The system as recited in claim 50, wherein the activity of the consistency coordinator is communicated to the caches.

10 52. The system as recited in claim 50, wherein the activity is communicated with heartbeat messages to the caches.